

IM3SHAPE: A maximum-likelihood galaxy shear measurement code for cosmic gravitational lensing

Joe Zuntz,^{1234*} Tomasz Kacprzak,¹ Lisa Voigt,¹ Michael Hirsch,¹ Barnaby Rowe¹
and Sarah Bridle¹²

¹*Department of Physics & Astronomy, University College London, Gower Street, London, WC1E 6BT*

²*Jodrell Bank Centre for Astrophysics, School of Physics and Astronomy, The University of Manchester, Manchester M13 9PL*

³*Astrophysics Group, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH*

⁴*Oxford Martin School, University of Oxford, Old Indian Institute, 34 Broad Street, Oxford OX1 3BD*

5 February 2013

ABSTRACT

We present and describe IM3SHAPE, a new publicly available galaxy shape measurement code for weak gravitational lensing shear. IM3SHAPE performs a maximum likelihood fit of a bulge-plus-disc galaxy model to noisy images, incorporating an applied point spread function. We detail challenges faced and choices made in its design and implementation, and then discuss various limitations that affect this and other maximum likelihood methods. We assess the bias arising from fitting an incorrect galaxy model using simple noise-free images and find that it should not be a concern for current cosmic shear surveys. We test IM3SHAPE on the GREAT08 Challenge image simulations, and meet the requirements for upcoming cosmic shear surveys in the case that the simulations are encompassed by the fitted model, using a simple correction for image noise bias. For the fiducial branch of GREAT08 we obtain a negligible additive shear bias and sub-two percent level multiplicative bias, which is suitable for analysis of current surveys. We fall short of the sub-percent level requirement for upcoming surveys, which we attribute to a combination of noise bias and the mis-match between our galaxy model and the model used in the GREAT08 simulations. We meet the requirements for current surveys across all branches of GREAT08, except those with small or high noise galaxies, which we would cut from our analysis. Using the GREAT08 metric we obtain a score of $Q=717$ for the usable branches, relative to the goal of $Q=1000$ for future experiments. The code is freely available from <https://bitbucket.org/joezuntz/im3shape>.

Key words: methods: statistical, methods: data analysis, techniques: image processing, cosmology: observations, gravitational lensing: weak, dark energy

1 INTRODUCTION

Weak gravitational lensing has the most potential for constraining the cosmological parameters that describe the dark universe (Albrecht et al. 2006; Peacock & Schneider 2006). As dark energy slows the growth of dark matter structures, these in turn affect the appearance of distant galaxies, distorting them slightly as gravitational fields deflect their light. This distortion is known as *cosmic shear*. Its measurement from individual galaxy images is crucial to the field of weak gravitational lensing, as the amount of distortion is directly related to the amount of intervening dark matter. However, the stretch induced by cosmic shear is typically on the order of a few percent, which renders accurate shape measurement an intricate but also intriguing problem.

Several upcoming and future observational surveys plan to

capitalise on the potential of cosmic shear for discovering the nature of dark energy. These include imminent ground-based projects (the Kilo-Degree Survey, Hyper Suprime-Cam¹ and the Dark Energy Survey: DES²), ground-based telescopes under construction (the Large Synoptic Survey Telescope: LSST³), and future space telescopes (*Euclid*⁴ and *WFIRST*⁵).

It is perhaps surprising that a major challenge for these experiments is simply measuring the shapes of the galaxy images. The difficulties arise because the galaxies involved are: small, less than the Point Spread Function (PSF) size in many cases; noisy, with a signal-to-noise ratio (S/N) of ~ 10 ; often irregular; and convolved

¹ <http://www.naoj.org/Projects/HSC/HSCProject.html>

² <http://www.darkenergysurvey.org>

³ <http://www.lsst.org>

⁴ <http://sci.esa.int/euclid>

⁵ <http://wfirst.gsfc.nasa.gov>

* E-mail: jaz@star.ucl.ac.uk

with a PSF which itself induces ellipticity in the galaxy that is typically greater than the gravitational shear we wish to measure. To fulfill the potential of cosmic shear these problems must be overcome, allowing precise and accurate measurements of the underlying gravitational shear.

1.1 Existing methods

The best known and most commonly-used method in weak lensing shape estimation remains that described by Kaiser et al. (1995), Luppino & Kaiser (1997) and Hoekstra et al. (1998), known as KSB or KSB+ (see also Schneider 2006 for an overview). KSB uses weighted quadrupole moments on galaxy images and PSF images to provide an estimator of gravitational shear. Many implementations of this approach, and derived methods adapted for space-based data (e.g. Rhodes et al. 2000, Schrabback et al. 2010), have been used to estimate gravitational shear over more than a decade. An important discovery of the Shear TEsting Programme (STEP) competitions was the sensitive dependence of the accuracy of KSB upon the details of its implementation (see Section 1.2, below; Heymans et al. 2006; Massey et al. 2007).

Recent modifications to KSB have sought to improve accuracy by adding higher-order terms to series approximations in the method (e.g. Okura & Futamase 2009; Viola et al. 2011; Melchior et al. 2011; Okura & Futamase 2012). More fundamentally different shape measurement approaches have also been proposed. A relatively early example was the ‘commutator’ method of Kaiser (2000); although this method has not found broad adoption in practice, the work contained a number of insights that were to inform the later development of the field.

Gaussianization, or Re-Gaussianization, approaches (e.g., REGLENS: see Mandelbaum et al. 2012, Section 4.3; Mandelbaum et al. 2005; Hirata & Seljak 2003) first convolve galaxy images with an additional kernel designed to make the final PSF as Gaussian, and isotropic, as possible. Such images then closely match many of the implicit assumptions in weighted quadrupole moment approaches such as KSB, so moments can then be used to extract shear estimates with reduced bias. Gaussianization methods remain competitive in terms of systematic bias (Mandelbaum et al. 2005, 2012) and have been used successfully in a number of studies in the Sloan Digital Sky Survey (e.g. Huff et al. 2011; Reyes et al. 2010; Mandelbaum et al. 2006; Hirata et al. 2004).

A family of methods known commonly as *shapelets* construct a model of PSF-convolved galaxy images as a sum of orthonormal Gauss-Laguerre polynomial functions (e.g. Bernstein & Jarvis 2002; Refregier & Bacon 2003; Massey & Refregier 2005; Kuijken 2006; Massey et al. 2007; Nakajima & Bernstein 2007). Effectively the ellipticity of derived shapelet models can then be used to provide an estimator of gravitational shear. Melchior et al. (2010) raised some conceptual concerns with the approach, and while similar (perhaps more drastic) conceptual problems with KSB did not affect its popularity there are fewer examples of shapelet shear estimation from real data in the literature (although see Velandier et al. 2011).

The number of proposed shape measurement methods is steadily growing. Interest in the problem has undoubtedly been stimulated by impending survey data of unprecedented statistical power, and by the blind analysis contests we describe in Section 1.2, below. Examples of innovative recent proposals the competitions described in the next section, and include Fourier Domain Null Testing (Bernstein 2010) and shear estimation via

lookup table (e.g., ‘MegaLUT’; see the review in the appendices of Kitching et al. 2012).

The final family of shape measurement methods are commonly known as model-fitting methods, and involve the comparison of simple parametric galaxy models (convolved by the PSF) to imaging data. The galaxy model ellipticity, among other parameters inferred in the model fit, make up the shear estimator.

The *lensfit* (Miller et al. 2007; Kitching et al. 2008; Miller et al. 2012) method seeks to perform Bayesian inference on galaxy model parameters and performed well in the GRavitational lEnsing Accuracy Teasting (GREAT08) Challenge: see Section 1.2, below; (Bridle et al. 2010). The more recent DeepZot method that also used a model-fitting algorithm proved the eventual winner of GREAT10 (Kitching et al. 2012).

The im2shape public code (Bridle et al. 2002)⁶ modelled both the galaxy and the PSF as a sum of Gaussians, inspired by Kuijken (1999). It performed MCMC sampling and took the mean of the of the ellipticity values of the samples as the shear estimate. It was used by Bardeau et al. (2005, 2007) to measure galaxy cluster masses.

1.2 Shear measurement contests

Researchers have come together to assess these many shear measurement methods in different regimes. The Shear TEsting Programme (STEP) was a joint blind analysis of simulated data by groups from within the shear measurement community (Heymans et al. 2006; Massey et al. 2007). It showed that the shear measurement problem is far from trivial, but demonstrated that the cosmological constraints set at that time were not limited by shear measurement accuracy. The first two GRavitational lEnsing Accuracy Testing (GREAT) Challenges posed the problem to computer scientists and provided sufficient simulations to assess whether methods were suitable for future surveys (Bridle et al. 2008; Kitching 2010). They showed that there are some regimes where existing methods work well enough, but others where further progress is necessary (Bridle et al. 2010; Kitching et al. 2012). The third challenge in the series, GREAT3,⁷ is currently in progress towards a challenge opening and simulation data release in 2013. In addition to a large simulation dataset, GREAT3 is leading the collaborative development of an open source extragalactic image simulation software toolkit,⁸ further advancing the STEP and GREAT programme of providing the community with a common reference for comparison and improvement of galaxy shape measurement methods.

1.3 IM3SHAPE

This paper presents a successor code to the im2shape code described above, called IM3SHAPE. It uses (by default) a sum of two Sersic profiles for the galaxy, similar to *lensfit* (Miller et al. 2007; Kitching et al. 2008; Miller et al. 2012), and a Moffat profile for the PSF. It calculates the maximum likelihood parameter values and corrects for noise bias (see, e.g., Refregier et al. 2012; Melchior & Viola 2012; Hirata et al. 2004; Bernstein & Jarvis 2002) using the calibration scheme described in Kacprzak et al. (2012).

⁶ <http://www.sarahbridle.net/im2shape/>

⁷ <http://great3challenge.info>

⁸ <https://github.com/GalSim-developers/GalSim>

Table 1. Shear measurement requirements on current, upcoming and future surveys. Requirements on the shear multiplicative and additive bias parameters m_i and c_i , and the GREAT08 Q value are shown for survey parameters: area A , number of galaxies per square arcminute n_{gal} and median redshift z_m .

Survey	A	n_{gal}	z_m	m_i	c_i	Q
Current	170	12	0.8	0.02	0.001	43
Upcoming	5000	12	0.8	0.004	0.0006	260
Future	2×10^4	35	0.9	0.001	0.0003	990

In Section 2 we discuss the requirements for shape measurement for current and future data sets. Then in Section 3 we present the IM3SHAPE code that we are releasing, the details of its operation and the choice of parameters it uses. Section 4 is concerned with noise-free simulations, to verify the code in ideal circumstances and to explore what biases the use of an incorrect model can introduce to the results. We present the results of tests of the code on the GREAT08 simulation set in Section 5, and conclude in Section 6. We give more details of the code implementation in Appendix A, of the termination criteria in Appendix B and overview the analysis process in Appendix C. Appendix D gives the equations and details used for analytic calculation of likelihood gradients with respect to the parameters.

2 REQUIREMENTS

Following Heymans et al. (2006), the accuracy of shear measurement methods are generally quantified in the literature in terms of the multiplicative error m and additive error c on the true shear γ^t , such that

$$\hat{\gamma}_i = (1 + m_i)\gamma_i^t + c_i \quad (1)$$

where $\hat{\gamma}_i$ is the shear estimate and the subscript i refers to the two shear components. It is assumed there is no cross contamination between, for example, $\hat{\gamma}_1$ and $\hat{\gamma}_2$.

The requirements on m and c for cosmic shear are computed in Amara & Réfrégier (2008) as a function of the survey area, depth and galaxy number density (see also Kitching et al. 2009). Amara & Réfrégier (2008) consider a two-point statistical analysis of the shear field, and the requirements are set so that the systematic error equals the statistical error, thus providing an upper limit on the allowed bias.

The GREAT08 (Bridle et al. 2008) challenge quantifies the accuracy of shear measurement methods using a single number called the *quality factor*, Q . The equation relating Q to m and c is given in Voigt & Bridle (2010). A related but different definition is used in GREAT10 (Kitching et al. 2012) but we stick to the GREAT08 definition here because we use the GREAT08 simulations.

In Table 1 we show the requirements on m , c and Q for three sets of survey parameters, chosen to represent current, upcoming and future cosmic shear surveys. Examples of surveys with requirements comparable to the ‘current’ set of requirements are the Canada-France-Hawaii Telescope Lensing Survey (CFHTLenS: Heymans et al. 2012), and likely the first results from KIDS, DES and HSC (de Jong et al. 2012; The Dark Energy Survey Collaboration 2005; Takada 2010). The upcoming survey requirements correspond roughly to the full DES and HSC surveys. The future survey requirements correspond to LSST and *Euclid* (e.g. Chang et al. 2012; Laureijs et al. 2011).

3 IM3SHAPE METHODOLOGY & CODE

The code we have developed and present in the remainder of this paper is called IM3SHAPE. It belongs to the family of model-fitting methods: it forward-fits a galaxy model to each data image it is supplied with and reports the parameters of the best fitting model, including the ellipticity components. For finding the best fit we use a maximum likelihood approach, i.e. we search for those model parameters that minimize χ^2 , the summed square difference between our generated model and the actual pixel data, weighted according to a user-supplied weight image. For details on the maximum likelihood estimation see Appendix A and D.

Throughout we assume that we have selected a square portion of the image, with *stamp_size* pixels on a side.

3.1 Code layout

IM3SHAPE is a modular C code, with a significant amount of Python glue code that enables the setting up of new models and their options automatically. Key processes like optimization and galaxy model generation routines are encapsulated in complex structures with simple interfaces, rather than complicated functions. This architecture makes it particularly easy to design new galaxy or star models to be fit to data. Since the process of fitting to a large number of postage stamps is inherently trivial to parallelize, the core of the code is all single-threaded.

3.2 Model choice

Using the flexible systems of IM3SHAPE we have experimented with fitting a variety of models, mainly to the images in the GREAT08 challenge.

We chose by default a two-component Sérsic, with a de Vaucouleurs bulge and an exponential disc, with the same centroid and ellipticities. In addition we simplify the fitting by fixing the ratio of scale radii, at unity.⁹

Within that general model we can choose different parameterizations for some components - these are discussed in detail in Appendix A.

Our complete set of parameters is shown in Table 2. The parameters marked as fixed were kept at default values during our standard analysis but can be varied easily if desired or to study model bias.

3.3 Model fitting

Generating the model image described in the previous Section is not trivial: we need to ensure that the model is at high enough resolution to capture sharp central galaxy features. This is discussed in Section 4.1.

Assuming images of the galaxy model can be accurately generated (see Section 4.1), maximizing the computed likelihood (i.e. minimizing the weighted, squared difference between the non-linear model and the data) is technically and computationally challenging. Reaching the current level of speed and accuracy in maximization within IM3SHAPE required considerable exploration of

⁹ This is similar but not identical to the model adopted in the GREAT08 simulations. IM3SHAPE should therefore work reasonably well on those images, and we note that this model was originally chosen to be a reasonable description of real galaxies.

Parameter	Meaning	Fixed
x_0	Horizontal centroid	
y_0	Vertical centroid	
e_1	x-y shear	
e_2	45° shear	
r_d	Disc half-light radius	
A_b	Bulge peak flux	
A_d	Disc peak flux	
R_r	Bulge-disc size ratio	✓
n_d	Disc Sérsic index	✓
n_b	Bulge Sérsic index	✓
Δe	Bulge-disc ellipticity	✓
$\Delta\theta$	Bulge-disc angle	✓

Table 2. The complete set of parameters used in the default galaxy model. Fixed parameters were not by default varied during parameter optimization.

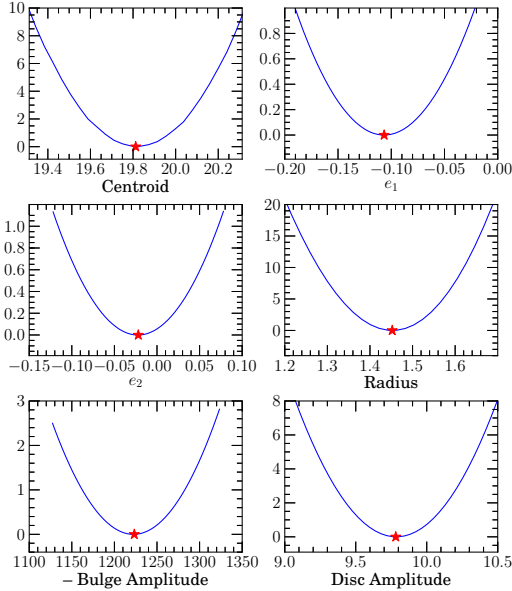


Figure 1. Slices in minus log-likelihood in six of the principal parameters (see section A) around the best fit point found for one galaxy by the likelihood maximizer, which is shown by a star.

maximization algorithms and their settings. We therefore give a brief overview of our findings in Appendix B, and an overview of the whole process of image generation and fitting in Appendix C.

With tuned parameter settings the maximizer, which uses a Levenberg-Marquardt algorithm, works well. Figure 1 shows example likelihoods from the first galaxy of set 3 in GREAT08. It can be seen that the point is indeed a good minimum (and further searches can confirm it to be a global one). The asymmetry in some of the parameters is part of the cause of the noise bias phenomenon (Kacprzak et al. 2012).

4 NOISE-FREE SIMULATIONS

Generating the model image described in the previous Section is not trivial; for example, we need to ensure that the model is at high

enough resolution to capture sharp central galaxy features without aliasing. We also need to understand how best-fitting models may be biased in cases where they cannot represent the data perfectly. In this Section we investigate the behaviour and performance of IM3SHAPE on noise-free simulations to find optimal parameter settings for model generation, and analyse in a simple way what potential biases can arise when the model is insufficiently flexible to describe the data.

4.1 Upsampling biases

The IM3SHAPE software uses the Discrete Fourier Transform (DFT) to render images of convolved galaxy profiles: this can be done with speed thanks to the Fast Fourier Transform (FFT) algorithm, freely available in highly efficient implementations (Frigo & Johnson 2005). However, representing the convolution of telescope PSFs with analytic, continuous, non-bandlimited functions such as the Sérsic profile requires careful numerical approximation, as can (if the PSF is also not bandlimited) the subsequent rendering of convolved profiles into pixellated images (see, e.g., Marks 2009). Decisions must be made about how and where to make approximations in the representation of these profiles, balancing precision against computational cost.

Upsampling, which is the more accurate rendering of model images using higher resolution (upsampled) intermediate images for internal operations such as convolution (see, e.g., Rowe et al. 2012), is a key tool to eliminate defects in modelling due to aliasing. We therefore investigate the requirements on the upsampling parameters used by the image generation code in IM3SHAPE. These three resolution parameters are:

- (1) U_1 , called *upsampling* in parameter files - the convolution between the galaxy model and the PSF model is performed on a grid of $(U_1 \times \text{stamp_size})^2$ pixels in size, i.e. each image pixel is divided into U_1 sub-pixels across. The flux within each sub-pixel is computed assuming the intensity is constant within it and equal to the intensity at the centre of the sub-pixel, for all pixels except the central N_u^2 pixels (see (2)).
- (2) N_u , called *n_pixels_to_upsample* in parameter files - within a central galaxy region of N_u^2 image sub-pixels the flux is integrated on a even finer grid within each sub-pixel (instead of assuming a constant intensity across each sub-pixel).
- (3) U_2 , called *n_central_pixel_upsampling* in parameter files - determines the number of integration intervals within a sub-pixel in the central image region; therefore the overall resolution in the central galaxy region is $U_1 \times U_2$ times the data image resolution.

The required values of these resolution parameters depend on the survey (see section 2). In figure 2 we plot the multiplicative bias as a function of the U_1 parameter for different values of N_u and with $U_2 = 7$. (The additive bias is negligible in comparison to the multiplicative bias because we use a circular PSF for this investigation.) Results are shown for a pure bulge and a pure disc with $R_{gp}/R_p = 1.4$ and 1.14, which are similar to the fiducial and small size branches used in GREAT08 (see Bridle et al. 2010, including for size definition). The bias is computed using a ring-test (Nakajima & Bernstein 2007) with 4 angles, $\{0, 45, 90, 135\}$ degrees.

The multiplicative bias is obtained by shearing each galaxy in the ring by an amount $\gamma_1 = 0.02$ along the x -axis and $\gamma_2 = 0.01$ along the line $y = x$. The PSF is a circular Moffat profile with

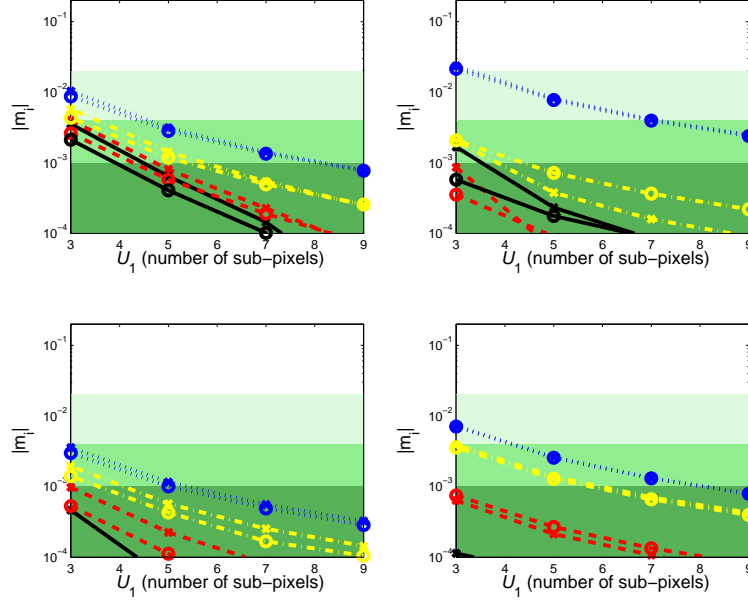


Figure 2. Resolution requirements: multiplicative bias (m_1 crosses, m_2 open circles) as a function of the number of sub-pixels U_1 used in the convolution between the galaxy model and the PSF model. Black solid, red dashed, yellow dot-dashed and blue dotted lines are for integration of the galaxy flux within 7x7, 5x5, 3x3 and 1x1 sub-pixels around the centre of the galaxy. The flux is integrated over a 7x7 grid within each sub-pixel. Top panels: small ($R_{gp}/R_p = 1.14$) galaxies. Bottom panels: fiducial size ($R_{gp}/R_p = 1.4$) galaxies. Left panels: de Vaucouleurs galaxies. Right panels: exponential galaxies. The shaded regions show the requirements on m_i given in Table 1. The upper edge of each shaded region (from bottom to top) shows the upper limit on the bias allowed for future, upcoming and current surveys respectively.

FWHM=2.85 pixels and shape parameter $\beta = 3.5$. The postage stamp size (*stamp_size* in parameter files) is 25 pixels. We find that the requirements on the resolution parameters do not change if we increase the stamp size to 45 pixels. The bulge-to-disk scale size ratio, R_r , is fixed to 1 during the fitting.

We find that for an upcoming survey (medium green shaded area in figure 2) we require $U_1 \geq 5$, $N_u \geq 3$, and $U_2 \geq 7$. For $U_2 < 7$ the results are unstable.

4.2 Model bias

Having established that the resolution parameters described above are high enough to remove any bias from aliasing effects, we now turn to the issue of *model* or *external* bias arising from fitting an incorrect model to the data. If a forward model approach could be perfectly implemented and the posterior distribution accurately characterized for each galaxy (rather than being reduced to a single point-estimate like the maximum likelihood) and combined appropriately, then this would be the only remaining source of bias for simple images like those in GREAT08. In practice, there may be other problems such as neighboring galaxies and image artefacts.

Very realistic galaxy models are almost certainly out of reach of forward methods for the foreseeable future - modelling spiral arms, for example, is hard enough for a single high-resolution image, let alone 10^8 noisy ones. We defer testing on extremely realistic models (for example using the SHERA simulation code of Mandelbaum et al. 2012, or Galsim), and instead apply a necessary but not sufficient test - fitting a simple but incorrect model using IM3SHAPE.

Again, we perform these tests without noise, to separate the issue of *noise* bias on maximum-likelihood estimates from *external* bias; of course, there may be an interaction between the two. As we

only touch upon these issues here, we refer the interested reader to Voigt & Bridle (2010) and Melchior et al. (2010) for a more extensive discussion.

Except where otherwise noted, our galaxy simulations in this section used $R_{gp}/R_p = 1.4$, equal component scale radii, and half the flux in each of the bulge and disc components $F_B/(F_B + F_D) = 0.5$. The PSF is a Moffat with $e_1 = e_2 = 0.03$, $\beta = 3$, and FWHM = 2.85.

4.2.1 Sérsic Indices

We first test a simple difference between the input (true) model and the fitted one - a different index of the Sérsic profile. Our input (true) model is generated with a single component, a Sérsic profile with index n_s , however, we fit our standard bulge plus disc model.

Figure 3 demonstrates this bias: it shows the pre- and post-PSF convolved simulated galaxy images, and the residuals between the simulated galaxy image and best-fit galaxy model, for two different values of the true Sérsic index. The images are noise-free and the correct PSF model has been used in the fits. The error on the true galaxy ellipticity is approximately 10^{-3} for $n_s = 1.7$, and 2×10^{-4} for $n_s \sim 3.5$.

The residual image shows different behaviour along the x and y directions, which is what causes the ellipticity to be mis-estimated. As discussed in Voigt & Bridle (2010), the galaxy size is worst estimated along the direction where the galaxy is smallest relative to the PSF. Therefore the ellipticity is incorrect.

Our model used in IM3SHAPE is correct when $n_s = 1$ or $n_s = 4$, since at these points the true model is a subset of the IM3SHAPE model. We generate simulations with our fiducial sets of galaxy parameters, except for the Sérsic index, which we vary. We apply a ring test to the simulations to assess the multiplicative

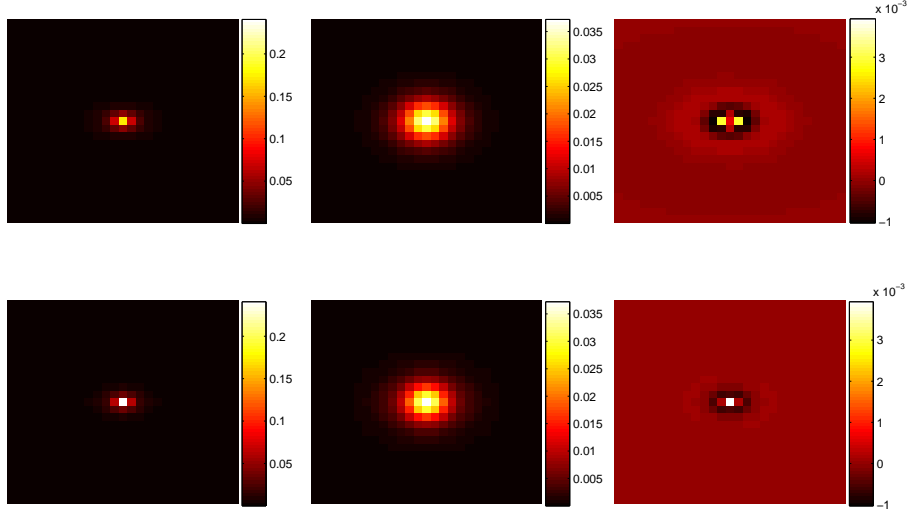


Figure 3. Examples of model bias: IM3SHAPE co-elliptical bulge ($n_s = 4.0$) plus disc ($n_s = 1.0$) model fits to simulated single-component Sérsic galaxies with $n_s = 1.7$ (top panels) and $n_s = 3.5$ (bottom panels). Postage stamp images are shown for the simulated galaxy (left-hand panels), the PSF-convolved simulated galaxy (middle panels), and the residuals between the simulated galaxy and the best-fit galaxy model from IM3SHAPE (right-hand panels). The simulated galaxy has ellipticity components $e_1 = 0.2$ and $e_2 = 0.0$ and the radius along the major axis is 1.84 pixels. The PSF is a circular Moffat profile with $\beta = 3.0$ and FWHM=2.85 pixels.

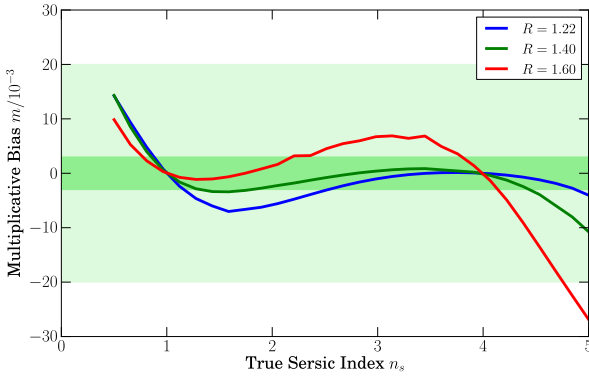


Figure 4. The mean multiplicative bias m as a function of true model Sérsic index when fit with a bulge+disc model, for three different image sizes as shown. The shaded bands are the current and upcoming survey target biases. Since the simulations were noise-free error bars are negligible.

m and additive c biases; the former is shown in figure 4, a figure for the latter is omitted as it is negligible in the noise-free case. As expected, the bias becomes zero when the true and fitted model coincide at $n_s = 1$ and $n_s = 4$.

Figure 4 suggests that this kind of model bias has the potential to be a problem for upcoming surveys but is only marginally relevant for current generation of experiments. Allowing for additional flexibility of the model could resolve this potential issue. Further tests are needed on more realistic galaxy images to assess the true impact on upcoming surveys.

4.2.2 Different bulge and disc ellipticity

In figure 5 the multiplicative bias arising from a slightly more realistic model difference is shown - a bulge and disc model with components that are not co-elliptical. This is of interest because it

is a difference between the IM3SHAPE model and the GREAT08 simulations. The input image bulge ellipticity¹⁰ e_b is fixed at 0.3 and the disc ellipticity varied so that $e_b = e_d + \Delta e$. The difference between the bulge and disc ellipticity, Δe , is varied between -0.25 and 0.25. Typically, bulges are rounder than discs, i.e $\Delta e < 0$. The model fit to the data is our standard two-component co-elliptical profile.

Once again a ring test is applied to assess the bias, and the noise is zero. The other galaxy parameters are set to the fiducial values described above, except that the different curves show different relative amplitudes of the bulge and disc components. The bias is zero when either Δe is small or the galaxy is completely bulge- or disc-dominated, since in that case the secondary component does not affect the fit at all. The effect is greatest when the components have comparable total fluxes, and this bias too is small for the current generation of experiments, if bulges are rounder than discs.

4.2.3 Incorrect Radius Ratio

As depicted in Table 2, the ratio between the half-light radii of the bulge and the disc is fixed by default in the code. This is partly to reduce the number of parameters (since fewer parameters generally means easier optimization), but also because this particular parameter causes problems with convergence in some regions of its range.

This is clearly a potential source of model bias - the ratio of these radii varies in real galaxies (when bulges and discs can be clearly identified). Furthermore, in GREAT08 this value is around $R_r \sim 1.5$. We again apply a ring test to noise-free simulations to check how severe this problem is. If we alter the ratio without changing the disc radius, then the bulge radius changes accordingly.

¹⁰ We define ellipticity as $e \equiv \sqrt{e_1^2 + e_2^2}$ with $e_{1/2}$ from Table 2 and Appendix A.

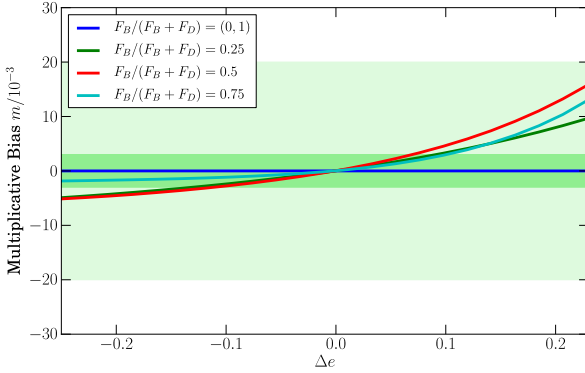


Figure 5. The mean multiplicative bias m as a function of true model Δe , the difference between bulge and disc ellipticities, when fit with a co-elliptical bulge+disc model. The shaded bands are target biases for current and upcoming experiments. Since the simulations were noise-free error bars are negligible. In real galaxies Δe is usually negative.

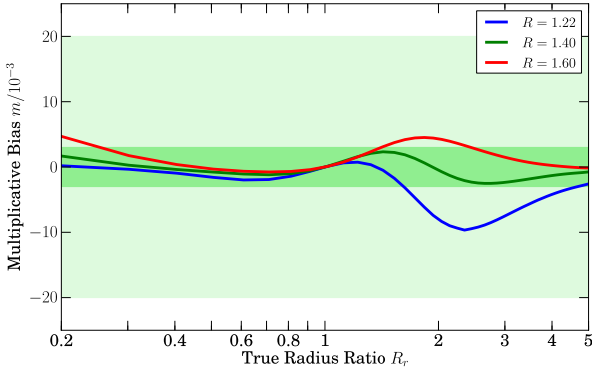


Figure 6. The mean multiplicative bias m as a function of the true ratio between the bulge and disc half-light radii for various galaxy sizes is shown. Note that the bias arises as the ratio has been fixed to one during the fitting with IM3SHAPE. Again, the simulations are noise-free so there are no error bars. The shaded bands are target biases for current and upcoming experiments.

To separate the effect of fitting small galaxies from this model bias we keep the ratio of the convolved image to PSF size, i.e. R_{gp}/R_p , fixed at three different values corresponding to GREAT08 sizes.

Figure 6 shows the results of this test. If real galaxies have smaller bulge scale radii than discs then the biases are mostly smaller than the requirements for upcoming surveys. If they are larger then the bias would be acceptable for current surveys but problematic for upcoming surveys, with a sign which depends on the galaxy size. For the GREAT08 fiducial branch ($R = 1.4$, true $R_r \sim 1.5$) we expect a slight over-estimate of the shear. This will combine with the slight underestimate expected from figure 5 for bulges which are rounder than discs, and may partially cancel out to bring results down to be compatible with the requirements for upcoming surveys. We recommend a program of testing on realistic images for calibration of upcoming surveys.

5 GREAT08

We tested IM3SHAPE on the Gravitational LEnsing Accuracy Testing 2008 (GREAT08) Challenge (Bridle et al. 2008, 2010) which has a sufficient number of galaxies to test methods at the accuracy required for future surveys (unlike STEP1 and STEP2). Furthermore, it features a constant shear across the images which greatly facilitates the computation of success metrics (unlike GREAT10) and therefore performance evaluation. The GREAT3 Challenge is still under construction (Mandelbaum, Rowe et al. in prep).

The GREAT08 Challenge includes two blind competitions: one with low noise, which has a signal-to-noise ratio¹¹ of 200 (‘LowNoise_Blind’) and one with more realistic signal-to-noise ratio around 20 (‘RealNoise_Blind’). In this section we show the results of IM3SHAPE on these simulations.

5.1 Low Noise Blind

The LowNoise_Blind challenge consists of 15 image files, each of which has 10,000 galaxy postage stamps of 39×39 pixels. All the galaxies are a mixture of bulge and disk components which are co-centered but not co-aligned. The misalignment angle is drawn from a Gaussian with a standard deviation of 20 degrees. We therefore expect a model bias to apply in these cases as described in section 4.2. The signal-to-noise ratio of 200 is large enough that noise bias described in Kacprzak et al. (2012) is expected to be negligible. The PSF and galaxy size parameters are described in detail in Bridle et al. (2010). In summary, there are 5 image files of each of 3 different galaxy sizes. This allows us to show how the success metrics depend on galaxy size.

The GREAT08 images are simulated at essentially infinite resolution by using photon shooting for the PSF and galaxy profiles (Bridle et al. 2010). IM3SHAPE makes its models at high resolution (see section 4.1). For this test we ran at two different values of this resolution, with an upsampling $U_1 = 5$ for a coarse run and $U_1 = 7$ for an improved, higher resolution run that took about twice as long to complete. We are interested to see both results on the small LowNoise_Blind dataset and would like to use the coarser resolution for the much larger RealNoise_Blind dataset.

In figure 7 we show the IM3SHAPE results on the LowNoise_Blind GREAT08 data. At high resolution the resulting Q -factor scores reach level $Q > 1000$ for all branches, a level at which the statistical noise in the challenge dominates - scores above this level are consistent with zero error. We reach this score in spite of a model bias: the the GREAT08 simulated galaxies have misaligned isophotes, whereas we fit co-elliptical ones, and the ratio of bulge to disc size ratio is different.

At the coarser resolution we are within the requirements for future surveys for all but the multiplicative bias on small galaxies, which is nonetheless acceptable for the current generation. We therefore use the coarser resolution in the remainder of this paper.

We also show the results submitted by others at the time of GREAT08. Significant developments have been made in other codes in the meantime, but we can conclude that at that time there were no methods as effective as IM3SHAPE in this high signal-to-noise regime. The closest competitor is the stacking method by AL, which has not yet been developed into a method that can be used on realistic data. We should however note that the model in the

¹¹ Here, we adopted the GREAT08 definition of signal-to-noise ratio defined by (A11) in Bridle et al. (2010), Appendix A, page 14.

IM3SHAPE code is very well matched to the GREAT08 simulations, in that they both use exponential disc plus de Vaucouleurs bulge galaxies. An investigation into this topic is beyond the scope of this work, and will be addressed by future challenges including GREAT3.

In running this analysis we determined which parameters in the code can be tweaked to minimize runtime while retaining high accuracy. The most important such parameters are those controlling the details of the LevMar optimization and convergence behavior. In subsequent work we have found that tuning these parameters is quite specific to the noise levels and other parameters of the galaxies in question, and we have not been able to draw any general conclusions about them.

5.2 Real Noise Blind

The GREAT08 RealNoise_Blind Challenge is considerably larger, with 2700 sets of 100×100 galaxy stamps. The 2700 sets are divided into 9 branches which span a range of simple galaxy models and PSF ellipticities and a range of realistic galaxy sizes and noise levels. The fiducial branch has a Signal-to-Noise a factor ten lower than in LowNoise_Blind but the same galaxy size as the central branch of LowNoise_Blind and the same PSF and galaxy type (co-centred but not co-elliptical bulge plus disc). Two branches explore the same two extrema in galaxy size as LowNoise_Blind; two branches vary the Signal-to-Noise by a factor either side of the fiducial value; two branches vary the PSF ellipticity by a factor of two and in direction; and two branches explore variants on the galaxy model: bulge or disc and bulge plus disc with an offset in the center. For more details on the challenge, see Bridle et al. (2008).

We tuned the maximization parameters described in Appendix B1 to improve the mean likelihood of all the galaxies in a single set; on the basis that if a set of minimisation parameters did not find the best likelihood then it must not have converged fully. We note that the mean ellipticity obtained for this single set was quite sensitive to the maximization parameters, and therefore we recommend users of the IM3SHAPE code should perform a similar tuning on new galaxy samples.

5.3 Noise Bias Calibration

At the noise levels in the RealNoise_Blind Challenge we do expect that we need to correct for the noise bias effect described in Refregier et al. (2012), Kacprzak et al. (2012) and references therein. We describe our approach to noise bias calibration in this subsection.

In Kacprzak et al. (2012) we ran IM3SHAPE on noisy simulations with a range of input parameters to find the behaviour of the noise bias with galaxy properties. We then attempted to remove the bias by matching the measured galaxy properties to the simulation input parameters to predict the expected bias, which we subtracted off. We ran this method on branch 3 of our GREAT08 analysis, which has the same galaxy model as IM3SHAPE, and found $Q = 166$, with $m = 3.0 \times 10^{-2}$ and $c = 7 \times 10^{-4}$ which is nearly satisfactory for current experiments but well outside our requirements for upcoming experiments.

This straightforward correction fails to do better because the observed properties are themselves noisy (and thus noise-biased), and so our bias estimate is itself biased. In practice, we have found it ineffective to fully calibrate galaxy ellipticity measurements one-by-one using their individual properties, and turn instead

to a population-based approach in which a complete class of galaxies is calibrated collectively.

The approach that we take is discussed in Appendix E. Briefly, we use the deep data from the low-noise part of the challenge to provide distributions of the galaxy properties, and then perform simulations to compute the mean multiplicative and additive biases for those properties. We then apply these mean biases to the shear estimates from each set.

To predict biases for a given parameter set we simulate galaxies at grid points with scale radii $r_d = \{1, 2.5, 5\}$ and bulge flux fractions $F = \{0, 0.5, 1\}$, and with ellipticity from $e_1 = -0.8$ to 0.8. At each grid point we fit a cubic polynomial to the bias $\hat{e} - e_{\text{true}}$ as a function of e_{true} , and store the coefficients. To get the bias for intermediate parameters we use a Gaussian radial basis function interpolator. Since the simulations are costly we simulate only at $S/N = 20$, and extrapolate to other values using the result from Refregier et al. (2012) and Kacprzak et al. (2012) that the bias scales as the inverse of the square of the signal-to-noise.

There are several factors about the structure of GREAT08 that are unrealistic in this context. Firstly, not every branch in the main sample has a corresponding low-noise version, whereas in real surveys it should be possible to match populations with more care in most cases. This would tend to make noise calibration harder on GREAT08. Secondly, the sizes of galaxies in GREAT08 are not drawn from a population - they have a single true value. This would tend to make calibration more effective than is realistic. We also note that the information needed to fully perform the calibration in this way was not public at the time of the challenge.

Where possible we match branches using the corresponding low-noise branch; where none exists we use the nearest approximation. See Appendix E for more details.

5.4 Scores

Q-factor scores and m and c values for the pre- and post-noise bias calibration results are shown in figure 8. Even before the calibration the bias on the $S/N = 40$ galaxies was small enough that we achieved a quality factor $Q > 1000$, with m and c small enough for upcoming surveys. At lower signal to noise, as we expect, the pre-calibration Q factor drops to below 100.

After the noise bias calibration procedure described above the Q factors and the c values reach the levels required for upcoming experiments for most of the branches of the challenge. The results are stable with PSF type. The m values are somewhat more variable but, depending on the exact branch, can reach the required levels. In particular the value is good in the b/d branch, where where we expect no model bias.

Our most encouraging result is the stability across the galaxy type branches. The only branch for which we use the correct model is the first galaxy type, where the simulation model is single-component bulge or disc. For all other branches we expect some model somewhat like the ones discussed in section 4.2, but as in that section the effect is not critical for current surveys.

The branches for small and noisy galaxies are clearly more problematic, and we would be forced to remove such galaxies from any current analysis. We believe the extrapolation fails badly in the high-noise case because of a slowdown in the magnitude of the bias with ellipticity, which does not follow the $b \sim (S/N)^{-2}$ relation that we used. Very noisy galaxies have much broader scatter in measured ellipticity, and when these values start to push up to the edge of the space, at $|e| = 1$, this acts as an $m < 0$ bias which

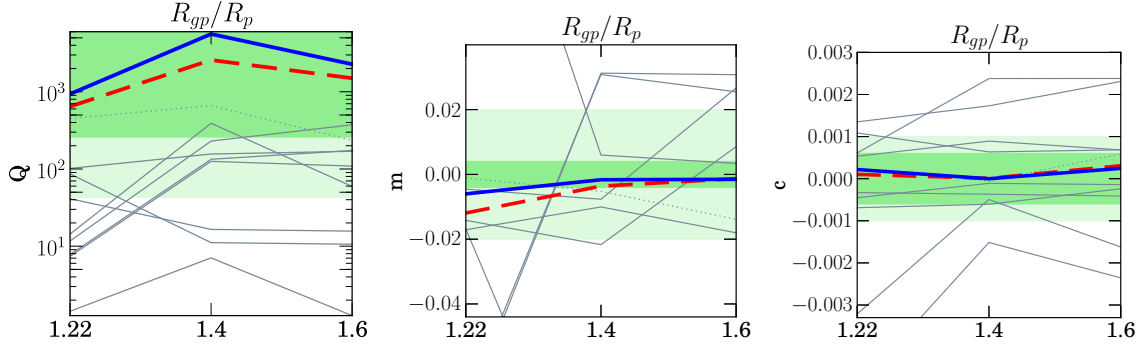


Figure 7. The Quality factor score Q , multiplicative bias m , and additive bias c (in the PSF direction) obtained in the LowNoise-Blind section of the GREAT08 challenge. The solid blue line is our main result, and the red dashed line is run at a lower resolution. The green bands are the requirements for current and upcoming surveys. The thin grey lines are other entrants at the time of the challenge, with dots indicating stacking methods.

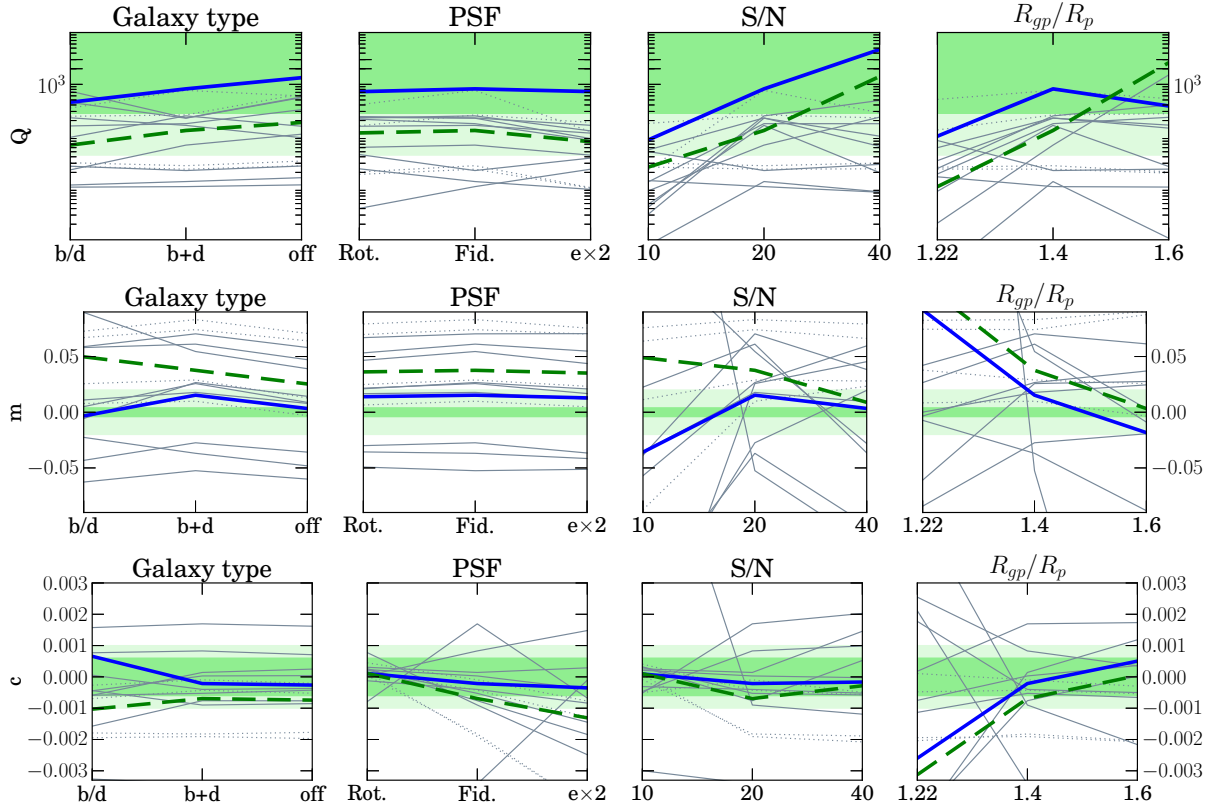


Figure 8. The Quality factor score Q , multiplicative bias m , and additive bias c obtained in the RealNoise-Blind section of the Great08 challenge. Each column shows scores with a different variation from the fiducial galaxy parameters. The thick solid blue line is our primary result, the IM3SHAPE scores following the matched-population noise bias calibration described in the text. The dashed green line is the score without any calibration. The green bands are the requirements for current and upcoming surveys. The thin grey lines are other entrants at the time of the challenge, with dots indicating stacking methods. The galaxy types are bulge or disc (50% of each; labelled b/d), co-elliptical bulge plus disc ($b + d$) and bulge plus disc with offset centroids (off). The PSF values were fiducial (fid), rotated by 90° , and double ellipticity ($e \times 2$).

partially counteracts the usual $m > 0$ one. Missing this effect can lead to extrapolated values over-correcting the bias.

6 CONCLUSIONS

We have presented a public galaxy shape fitting code, IM3SHAPE, for weak lensing shear measurement. Our code efficiently finds the

maximum likelihood fit of our model to any given galaxy, and does so extremely well: we find little sign of any bias due to poor minimization for simple simulated galaxies, following careful tuning of the model image resolution settings and other parameters. The code is modular and easily extensible, and we would welcome com-

munity contribution to any part of it¹². The technical problems in generating models and fitting them to galaxies are extensive, with particular problems of resolution, parameterization, and non-linear numerical optimization. We discuss our solutions to these problems in the appendices to this paper.

We find that we need to generate model images at five times the data resolution to reach our desired accuracy for upcoming surveys. We also find that to accurately capture small scale features in de Vaucouleurs cores we must sample a small central region even more finely - at 35 times the data resolution.

Our code fits a simple bulge plus disc galaxy model to galaxy images, and we have made an exploration of model bias, which can arise when true galaxies fail to be as neatly analytic in form as our model ones. We find that reasonable deviations between the simple true and assumed models cause only minor bias, not significant for current data and early results from the upcoming generation of surveys. The bias will, however, attain greater significance as upcoming surveys reach their full statistical potential. Our tests were rather limited in that we considered only simple differences from our canonical model - deviations in Sersic index, and differences between the size and ellipticity of the two model components. We look forward to performing more realistic tests with the GalSim tool set in future work (Mandelbaum, Rowe et al. in prep).

We have demonstrated the utility of our method and code by testing on the GRavitational lEnsing Accuracy Testing 2008 (GREAT08) Challenge in both low and realistic noise regimes. We find extremely good results for the low noise simulations, with multiplicative and additive shear measurement biases well within the requirements for upcoming (Stage III) surveys, for the medium sized and large galaxies. For practical reasons we chose to compromise on the resolution we used in the fits and thus our results for the smallest galaxies are only sufficiently good for current surveys.

In Kacprzak et al. (2012) we calculated the degradation of shear measurements due to noise on images and discussed correction schemes. We do not expect this to be significant at the signal-to-noise of the LowNoise simulations ($S/N = 200$). We expect the main difficulty to arise from resolution and the differences between our fitted model and the GREAT08 galaxy models. Both our galaxy models and the GREAT08 low noise simulations assume galaxies are made of a de Vaucouleurs bulge plus an exponential disc. The difference is that we assume these two components are co-centric and co-elliptical with the same size ratio, whereas in GREAT08 low noise the two components have different ellipticities, are not co-aligned and have a bulge scale radius which is approximately 50% larger than the bulge scale radius. We find that there is little resulting problem arising from model bias alone, as expected from our simpler investigations.

We ran IM3SHAPE on the 27 million galaxies in the realistic noise part of the GREAT08 Challenge, which we chose over GREAT10 because it is simpler and the variable shear fields used in the latter were not relevant to this method.

At this signal-to-noise ($S/N \sim 20$) we do expect significant noise bias due to pixel noise in the images (Kacprzak et al. 2012), leading to a multiplicative shear measurement error of a few per cent. We found that the correction scheme tested in (Kacprzak et al. 2012) does not work well on GREAT08 due to the bias-on-bias problem described in that work. We therefore deviated from the methods discussed there, which corrected each galaxy individually, in favour of a global correction using the knowledge of galaxy

properties measured in the corresponding low noise sample. The leading order correction to shear power spectrum measurements is expected to simply depend on the mean multiplicative bias of a population, which can be corrected globally. Since the noise bias depends on galaxy properties like size, and galaxies with similar properties can be clustered, then it is possible this scheme may not be satisfactory for high precision corrections.

We scale the noise bias correction from the fiducial value according to signal-to-noise, as expected in the analytic calculations of Refregier et al. (2012). This works well for the lower noise images but less well for those at higher noise levels as also noted in (Kacprzak et al. 2012). As stated in that work we are not satisfied with the performance of our optimiser in this highest noise regime and would cut such images from cosmic shear analyses at this point, or calibrate using a more extensive suite of population matched simulations. A similar concern applies to the smallest galaxies.

For all other branches of GREAT08 we attain additive and multiplicative biases which are within the requirements for current surveys, and have quality factor metrics which are as good or better than other non-stacking methods at the time of the GREAT08 Challenge.

Further work appears to be required before IM3SHAPE can be applied to the final years of upcoming Stage III cosmic shear surveys. We speculate that our residual shear measurement bias is due to the combined effects of using the incorrect galaxy model in the fit, and the noise on the images. Ideally an improved shear measurement method could be found that did not suffer from these types of bias. Alternatively the calibration scheme needs to be tailored more accurately to the analysis in question, for example by using a suite of low noise images with known shears.

To be fully applicable to real data there are a number of extensions required to the code, for example (i) the ability to handle multiple exposures of the same patch of sky (ii) the capacity to deal with or flag neighboring objects contaminating the shear measurement.

Furthermore, we have not investigated here the effect of postage stamp size or the effect of a weight map cutting out parts of the galaxy image. These developments are ongoing but testing them is beyond the scope of the present work.

ACKNOWLEDGMENTS

We thank Steve Gull, Jean-Paul Kneib, Eduardo Cypriano, Phil Marshall, John Bridle, Patrick Hudelot, Sebastien Bardeau, David Hogg, Alexandre Refregier, Adam Amara, Filipe Abdalla, Niall Maccrann, Emmanuel Bertin, Jim Bosch, Mike Jarvis, Fabrizio Sidoli, Gary Bernstein, Catherine Heymans and Tom Kitching for helpful discussions, and Dugan Witherick for patient and extensive computational help.

JZ, TK, SB, BR, LV and MH acknowledge support from European Research Council in the form of a Starting Grant with number 240672.

This paper has been typeset from a \TeX / \LaTeX file prepared by the author.

REFERENCES

Amara A., Réfrégier A., 2008, MNRAS, 391, 228

¹² <http://bitbucket.org/joezuntz/im3shape>

Bardeau S., Kneib J.-P., Czoske O., Soucail G., Smail I., Ebeling H., Smith G. P., 2005, *A&A*, 434, 433

Bardeau S., Soucail G., Kneib J.-P., Czoske O., Ebeling H., Hudelot P., Smail I., Smith G. P., 2007, *A&A*, 470, 449

Bernstein G. M., 2010, *MNRAS*, 406, 2793

Bernstein G. M., Jarvis M., 2002, *Astronomical Journal*, 123, 583

Bridle S. et al., 2010, *MNRAS*, 405, 2044

Bridle S., et al., 2008, *ArXiv* 0802.1214

Bridle S., et al., 2010, *MNRAS*, 405, 2044

Bridle S. et al., 2008, *ArXiv* 0802.1214

Bridle S. L., Kneib J.-P., Bardeau S., Gull S. F., 2002, in Natarajan P., ed., *The Shapes of Galaxies and their Dark Halos*. pp 38–46

Chang C. et al., 2012, *ArXiv e-prints: astro-ph/1206.1378*

de Jong J. T. A., Verdoes Kleijn G. A., Kuijken K. H., Valentijn E. A., 2012, *Experimental Astronomy*, p. 34

Frigo M., Johnson S. G., 2005, *Proceedings of the IEEE*, 93, 216

Heymans C., et al., 2006, *MNRAS*, 368, 1323

Heymans C. et al., 2012, *MNRAS*, 427, 146

Hirata C., Seljak U., 2003, *MNRAS*, 343, 459

Hirata C. M. et al., 2004, *MNRAS*, 353, 529

Hoekstra H., Franx M., Kuijken K., Squires G., 1998, *Ap. J.*, 504, 636

Huff E. M., Eifler T., Hirata C. M., Mandelbaum R., Schlegel D., Seljak U., 2011, *ArXiv astro-ph/1112.3143*

Kacprzak T., Zuntz J., Rowe B., Bridle S., Refregier A., Amara A., Voigt L., Hirsch M., 2012, *ArXiv* 1203.5049

Kaiser N., 2000, *Ap. J.*, 537, 555

Kaiser N., Squires G., Broadhurst T., 1995, *Ap. J.*, 449, 460

Kitching T. D., Amara A., Abdalla F. B., Joachimi B., Refregier A., 2009, *MNRAS*, 399, 2107

Kitching T. D., et al., 2012, *MNRAS*, 423, 3163

Kitching T. D., Miller L., Heymans C. E., van Waerbeke L., Heavens A. F., 2008, *MNRAS*, 390, 149

Kitching T. e. a., 2010, *ArXiv* 1009.0779

Kuijken K., 1999, *A&A*, 352, 355

Kuijken K., 2006, *A&A*, 456, 827

Laureijs R. et al., 2011, *ArXiv e-prints: astro-ph/1110.3193*

Luppino G. A., Kaiser N., 1997, *Ap. J.*, 475, 20

Mandelbaum R., Hirata C. M., Broderick T., Seljak U., Brinkmann J., 2006, *MNRAS*, 370, 1008

Mandelbaum R., Hirata C. M., Leauthaud A., Massey R. J., Rhodes J., 2012, *MNRAS*, 420, 1518

Mandelbaum R. et al., 2005, *MNRAS*, 361, 1287

Marks R. J., 2009, *Handbook of Fourier Analysis & Its Applications*, 1st edn. Oxford University Press

Massey R., et al., 2007, *MNRAS*, 376, 13

Massey R., Refregier A., 2005, *MNRAS*, 363, 197

Massey R., Rowe B., Refregier A., Bacon D. J., Bergé J., 2007, *MNRAS*, 380, 229

Melchior P., Böhnert A., Lombardi M., Bartelmann M., 2010, *A&A*, 510, A75+

Melchior P., Viola M., 2012, *MNRAS*, 424, 2757

Melchior P., Viola M., Schäfer B. M., Bartelmann M., 2011, *MNRAS*, 412, 1552

Miller L. et al., 2012, *ArXiv e-prints: astro-ph/1210.8201*

Miller L., Kitching T. D., Heymans C., Heavens A. F., van Waerbeke L., 2007, *MNRAS*, 382, 315

Nakajima R., Bernstein G., 2007, *Astronomical Journal*, 133, 1763

Okura Y., Futamase T., 2009, *Ap. J.*, 699, 143

Okura Y., Futamase T., 2012, *Ap. J.*, 748, 112

Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 1992. Cambridge: Cambridge Univ. Press

Refregier A., Bacon D., 2003, *MNRAS*, 338, 48

Refregier A., Kacprzak T., Amara A., Bridle S., Rowe B., 2012, *ArXiv* 1203.5050

Reyes R., Mandelbaum R., Seljak U., Baldauf T., Gunn J. E., Lombriser L., Smith R. E., 2010, *Nature*, 464, 256

Rhodes J., Refregier A., Groth E. J., 2000, *Ap. J.*, 536, 79

Rowe B., Bacon D., Massey R., Heymans C., Haeussler B., Taylor A., Rhodes J., Mellier Y., 2012, *ArXiv e-prints: astro-ph/1211.0966*

Schneider P., 2006, in Meylan G., Jetzer P., North P., Schneider P., Kochanek C. S., Wambsganss J., eds, *Saas-Fee Advanced Course 33: Gravitational Lensing: Strong, Weak and Micro*. pp 269–451

Schraback T., et al., 2010, *A&A*, 516, A63

Takada M., 2010, in Kawai N., Nagataki S., eds, *American Institute of Physics Conference Series Vol. 1279*, American Institute of Physics Conference Series. pp 120–127

The Dark Energy Survey Collaboration 2005, *ArXiv Astrophysics e-prints*

Velander M., Kuijken K., Schraback T., 2011, *MNRAS*, 412, 2665

Viola M., Melchior P., Bartelmann M., 2011, *MNRAS*, 410, 2156

Voigt L. M., Bridle S. L., 2010, *MNRAS*, 404, 458

APPENDIX A: MODELLING AND PARAMETERIZATION

A1 Choice of ellipticity parameters

The performance of all the optimizers we have tried is heavily dependent on the particular choice of parameterization. In particular, the choice of how to parameterize the shear components is particularly important since they are our target. We define ellipticity to match the effect of shear on images: $e = (a - b)/(a + b)$. The two obvious parameterizations based on this, (e_1, e_2) and (e, θ) , were both found to be problematic.

Using the two separate ellipticities e_1 and e_2 as parameters creates a non-linear boundary at $e_1^2 + e_2^2 = 1$, the edge of the unit circle.

While most optimization methods cleanly handle linear boundaries (for example simple limits on parameters), non-linear ones are less well served. We tried simply assigning an extremely low likelihood to any proposed sample beyond this boundary, but fast optimization methods use likelihood differences to determine step sizes, so they typically behave wildly if large negative values are used to signal edges. Since the images of model galaxies diverge as $|e| \rightarrow 1$ simply making the models of these galaxies is also prone to numerical problems.

Using the polar coordinates (e, θ) as parameters works well at high ellipticity where we use a simple parameter limit on e , but fails at lower ellipticities. As $e \rightarrow 0$ the likelihood becomes a plateau in θ since the coordinates are degenerate. Optimizers spent an extremely long time wandering around the circle.

We found a reliable parameter set which is a variant of the first parameterisation, (e_1, e_2) . We use (e_1, e_2) as the parameters varied by the optimizer, but map them to a disc wherever they stray beyond it. We mirror the parameter space radially about some e just below unity (we found 0.95 worked well), so that as the input e increases above the maximum, the value given to the minimiser first decreases and eventually reaches zero, and turns negative so that

the signs of the output e_1 and e_2 are flipped. Although this mirror structure causes repetition in the likelihood surface and multiple modes we have not found this to be an issue in practice.

A1.1 Other parameters

There are various radius parameters we could choose for each of the components. The choice is particularly important for the sharp bulge profiles, as their full-width half-maxima are extremely small and so they are not suitable parameters.

A better choice that works for both components is the radius which encloses half the total flux of the complete profile. Rather than using two radius parameters, we use the disc radius and the ratio between the radius of the bulge and disc. We find (see section 4.2) that no large model bias is generated by fixing the latter parameter provided it is approximately correct. We therefore use a fixed value for any given data set. For our GREAT08 runs this value was arbitrarily set to unity, to avoid trying to match too closely to the GREAT08 truth values.

A likelihood plateau occurs for small objects like stars, where the shears are irrelevant. In these regimes the fitted radii of the objects tended to fail to converge as they varied around zero. We solved this problem by setting a small transition radius $R = 0.4$ sub-pixels and a replacing the radius r with $\frac{1}{2}(R + r^2/R)$ when $|r| < R$ and $|r|$ otherwise. The minimum value radius actually used is then $R/2 = 0.2$ sub-pixels.

A1.2 Marginalization

We have experimented with marginalization over the amplitude components of the models. Since the model is linear in these parameters their likelihood through a given slice in the other parameters is Gaussian and we can analytically marginalize over them. We have found this to be slightly problematic in general - the marginalization is not particularly numerically stable and since we have two components we need to perform the PSF convolution twice if we sum them after marginalization. These two effects meant it was in fact slower to use this trick, provided that we have a reasonable initial estimate of the ellipticity, particularly as our gradient-based minimizer finds the linear parameters very simple to optimize.

If we move the modelling into Fourier space we could further marginalize over the centroid parameters x_0 and y_0 of the model - combining these dimensionality reductions could well yield a significant speed up. The downside of this is decreased flexibility in the modelling. We defer further discussion of this issue.

A1.3 Point spread function (PSF)

IM3SHAPE has the general facility to use an arbitrary image for PSF, but by default we read a catalog of either Moffat or Airy function parameters and use that to build a PSF image. It is built at the same upsampled resolution as the image itself, and transformed and stored in Fourier space for faster convolution with trial images.

When these images are made the convolution is done using the FFTW package Frigo & Johnson (2005).

A1.4 Fourier Sampling Kernel

Our model for the pixels in an image is really a constant value across the surface of the pixel. But as far as a Discrete Fourier transform is concerned a 2D array represents a grid of δ -function

samples from the image. To account for this difference our Fourier-space model of the image needs to include convolution with a square top-hat function the size of a pixel - this is true even when simulating at the same resolution as the data. Since we are also convolving with a PSF we can fold together these convolutions into a single Fourier-space product. The Fourier transform of a 2D top-hat is given by:

$$F_{ij} = \text{sinc}\left(\frac{i\pi u}{2n_x}\right) \text{sinc}\left(\frac{j\pi u}{2n_y}\right), \quad (\text{A1})$$

where the image dimensions are $(n_x \times n_y)$ and the upsampling factor is u .

A2 Resolution

The model image in IM3SHAPE is made at a resolution several times higher than the data resolution, and the convolution with the PSF also happens at this high resolution. Only then is the model downsampled to compare to the data. This procedure ensures that small and sharp features in the galaxy can be captured.

We find, though, that the central regions of galaxy bulges with the standard de Vaucouleurs profile are so sharp that an even higher resolution is needed to accurately model them. Rather than generate the whole image at this extreme resolution we only use it for the central pixels where it is relevant.

There are therefore three different resolution parameters to consider: the upsampling (the ratio of the model to image pixel size in the outer image region), the size of the central higher resolution region, and the resolution increase in this central region (compared to the outer region). The values that we used for these parameters are discussed in more detail in section 4.1.

APPENDIX B: NUMERICAL OPTIMIZATION

The optimization problem of finding the best model parameters can be tackled with a choice of algorithms; we have found that the best performing is the Levenberg-Marquadt (LevMar) method. This method uses more information than methods (such as Powell's method or Nelder-Mead) which operate only on the total likelihood as a function of the input parameters: it uses the complete mismatch map, (i.e. the χ^2 per pixel) for each parameter set considered.

The LevMar method uses gradients (analytic or numerical) to optimize; it is therefore not possible to signal a failure of a particular parameter set (which can occur when parameters are near their limits) using NaNs or other flag values. While we do try to put in place a parameterization that prevents this from happening (see below), we also signal the failure by setting the χ^2 for the failing model to 2 per pixel, a value large enough to be rejected but hopefully not hugely problematic.

B1 Termination criteria

The LevMar method depends on a number of manually set thresholds that determine when the numerical optimization should terminate. In particular, it stops and reports the parameters of the best model fit when at least one of the following criteria is met:

- One of the partial derivatives of the objective function (χ^2) with respect to the model parameters falls below a threshold ε_1 ,

i.e. $\max |\frac{\partial \chi^2}{\partial p}| < \epsilon_1$, where p denotes one of the model parameters. Note that this termination criterion can be sensitive to the changes of image pixel scaling and changes in the signal-to-noise ratio (SNR).

- The relative change of the norm of the estimated parameter vector is smaller than ϵ_2 , i.e. $\sum_p |p^{i-1} - p^i|^2 < \epsilon_2^2 \cdot \sum_p |p^{i-1}|^2$, where i is the current iteration index. This termination criterion can be sensitive to individual parameter scaling.

- The value of χ^2 falls below ϵ_3 ; again, this termination criterion is sensitive to image scaling and the S/N ratio.

- The relative change of χ^2 falls below ϵ_4 , i.e. $|(\chi^2)^{i-1} - (\chi^2)^i|/(\chi^2)^{i-1} < \epsilon_4/n$, where $(\chi^2)^i$ is the objective value at the current iteration and n is the total number of pixels. This termination criterion was implemented by the authors and is not included in the LevMar implementation we are using. It can be sensitive to the step size that LevMar is taking and thus will be affected by the `levmar_init_mu` parameter - the initial damping of the normal equations solved in the Levenberg - Marquadt algorithm.

- The maximum number of iterations is reached (the minimization has failed).

APPENDIX C: ANALYSIS OVERVIEW

Before analyzing a collection of galaxy images we:

- Read a collection of options specifying the model to be used, parameter ranges, and various options from a parameter file
- Build a model object encapsulating the model we will fit to the galaxies, including its likelihood and prior functions

Before analyzing each individual image we then:

- Generate the point-spread image based on PSF parameters; compute its FT.
- Generate the sinc function for the specified upsampling (the FT of a top-hat).
- Multiply these two Fourier space quantities together to generate the Fourier kernel.
- Generate a weight map based on the image noise levels.
- Compute a starting estimate for the galaxy properties from defaults and weighted quadrupole moments.

We then pass the model likelihood function to a minimizer object which runs the chosen algorithm. At each evaluation of the posterior, for our standard model, we:

- Check if the parameters are outside the desired ranges and give up early if not.
- Build high-resolution (up-sampled) real-space models of each component of the galaxy (e.g. bulge and disc).
- Replace the central few pixels of each model with summed up pixels from even higher resolution, since this is the region where the image changes quickly.
- Sum and convolve the components with the precomputed kernel.
- Down-sample the image to the data resolution.
- Compute the likelihood from the weight map and model image.
- Save the model image for use in the optimizer and residual analysis.

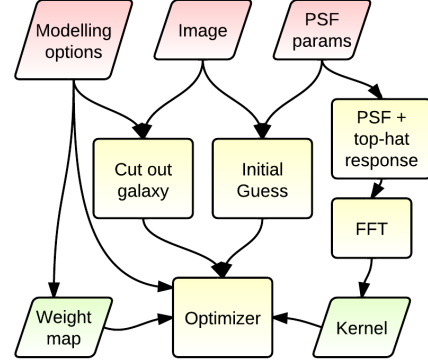


Figure C1. Actions in the IM3SHAPEcode run to prepare for the analysis of a single galaxy. Inputs, which can be collectively specified for a collection of galaxies, are upper red parallelograms, stored data are lower green parallelograms, and calculations are yellow rectangles. The preparation time is negligible compared to the time of the optimization itself.

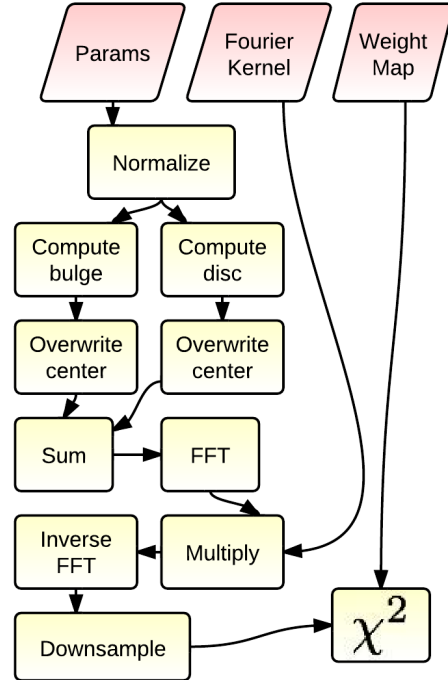


Figure C2. The calculation performed whenever the optimization code provides a set of parameters for which to calculate the image and likelihood. The principal parameterization used in this work and the normalization applied to it are described in section A. The “overwrite center” process is described in more detail in section 4.1.

APPENDIX D: DETAILS ON THE MAXIMUM LIKELIHOOD ESTIMATION

The maximization problem outlined in section B is based on the following model:

$$\mathbf{y} = \mathbf{D}(\mathbf{K} * \mathbf{x}(\mathbf{p})) + \mathbf{n}, \quad (\text{D1})$$

where \mathbf{y} denotes the observed galaxy image, $\mathbf{K} = \mathbf{f} * \mathbf{\Pi}$, the convolution of the PSF \mathbf{f} and a square top-hat pixel window function $\mathbf{\Pi}$ (see section A1.4), $\mathbf{x}(\mathbf{p})$ the (higher resolution) galaxy model with

parameter vector \mathbf{p} , and $D(\dots)$ the downsampling operator, which averages into lower resolution pixels. As detailed in section 3 we assume a two-component Sérsic galaxy model with the parameters $\mathbf{p} = (x_0, y_0, e_1, e_2, r_d, A_b, A_d)$ (see Table 2 for a detailed parameter description)

$$\mathbf{x}(\mathbf{p}) = \sum_{k=b,d} A_k \exp \left\{ -\mathbf{R}(x_0, y_0, e_1, e_2, r_d)^{1/2n_k} \right\}, \quad (\text{D2})$$

where $n_k = 1$ for the disc component and 4 for the bulges, and \mathbf{R} denotes the sheared squared distance

$$R_i = \mathbf{d}_i^T \mathbf{C} \mathbf{d}_i \quad (\text{D3})$$

with the centered coordinate vector $\mathbf{d}_i = (x_i - x_0, y_i - y_0)$ and the covariance matrix

$$\mathbf{C} = \frac{1}{r_d(1 - |e|)} \begin{pmatrix} 1 + |e| - 2e_1 & -2e_2 \\ -2e_2 & 1 + |e| + 2e_1 \end{pmatrix}. \quad (\text{D4})$$

For additive Gaussian noise, i.e. $\mathbf{n} \propto \mathcal{N}(\mathbf{0}, \mathbf{\Sigma})$ maximizing the likelihood $P(\mathbf{p}|\mathbf{y}, \mathbf{K}, \mathbf{x})$ is equivalent to minimizing

$$\chi^2(\mathbf{p}) = \frac{1}{2} \sum_i r_i(\mathbf{p})^2 \quad (\text{D5})$$

with respect to the model parameters \mathbf{p} , where \mathbf{r} denotes the residual image between the observed galaxy image \mathbf{y} and the downsampled, PSF convolved galaxy model image $\mathbf{x}(\mathbf{p})$, i.e.

$$\mathbf{r}(\mathbf{p}) = \frac{1}{\sigma} [\mathbf{y} - D(\mathbf{K} * \mathbf{x}(\mathbf{p}))]. \quad (\text{D6})$$

For simplicity we assume \mathbf{n} to be uncorrelated white noise and so consider only diagonal terms of the covariance matrix, i.e. $\sigma = \text{diag } \mathbf{\Sigma}$; we neglect any correlation in noise of adjacent pixels.

For minimizing (D5), IM3SHAPE uses the free software package LEVMAR¹³, a C/C++ implementation of the Levenberg-Marquardt (LM) algorithm, a standard technique for solving non-linear least-squares problems (Press et al. 1992). The LM method is based on a linear approximation of \mathbf{r} in the neighbourhood of \mathbf{p} , i.e.

$$\mathbf{r}(\mathbf{p} + \Delta\mathbf{p}) \approx \mathbf{r}(\mathbf{p}) + \mathbf{J}(\mathbf{p})\Delta\mathbf{p} + O(\|\Delta\mathbf{p}\|^2), \quad (\text{D7})$$

where $\Delta\mathbf{p}$ is a small perturbation of the parameter vector \mathbf{p} and \mathbf{J} denotes the Jacobian matrix of \mathbf{r} :

$$J_{ij}(\mathbf{p}) = \left. \frac{\partial r_i}{\partial p_j} \right|_{\mathbf{p}}. \quad (\text{D8})$$

The Jacobian (D8) can be numerically approximated, but to speed convergence of the minimization procedure it is generally recommended to compute it analytically. For our model (D1) the Jacobian reads with respect to the model parameters $p \in \{x_0, y_0, e_1, e_2, r_d\}$

$$\frac{\partial \mathbf{r}}{\partial p} = \frac{1}{\sigma} D \left(\mathbf{K} * \sum_{k=b,d} \left[\frac{A_k}{2n_k} \exp \left\{ -\mathbf{R}^{1/2n_k} \right\} \cdot \mathbf{R}^{1/2n_k - 1} \cdot \frac{\partial \mathbf{R}}{\partial p} \right] \right) \quad (\text{D9})$$

where \cdot denotes element-wise multiplication and the partial derivatives $\partial \mathbf{R} / \partial p$ are readily computed from (D3) and (D4). With respect to the bulge and disc amplitudes $A_{k \in \{b,d\}}$ one obtains

$$\frac{\partial \mathbf{r}}{\partial A_k} = -\frac{1}{\sigma} D \left(\mathbf{K} * \exp \left\{ -\mathbf{R}^{1/2n_k} \right\} \right). \quad (\text{D10})$$

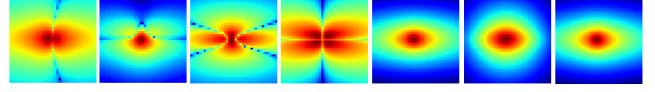


Figure D1. Components of Jacobian matrix with respect to its model parameters (from left to right): $x_0, y_0, e_1, e_2, r_d, A_b, A_d$.

Note that for clarity and the ease of exposition we omitted the image upsampling within a central image region as described in section 4.1 in the derivation above, i.e. we considered the special case where *n_central_pixel_upsampling* is set to one. The generalization to *n_central_pixel_upsampling* > 1 is straightforward since it only involves an additional summation over all *n_central_pixel_upsampling*² sub-pixels to yield the flux $[\mathbf{x}(\mathbf{p})]_{ij}$ for each pixel (i, j) within the central image region.

APPENDIX E: NOISE BIAS CALIBRATION USING DEEP DATA

Noise bias is the difference between the expectation of the maximum-likelihood result found by a model fitter like IM3SHAPE, which is a biased estimator, and the true underlying value, coming from the non-linear mapping between parameters and image. Its typical value is a few percent, and we need to remove it to reach our target accuracy levels.

The size of the bias depends sensitively on the true parameters of the galaxy, and if these were known we could remove noise bias completely. However, we have access only to noisy estimates of these parameters and therefore our estimate of the noise bias is itself noisy, and biased. We can think of this as a *bias-on-bias* problem, and we find that failing to account for it means we significantly miss our targets.

We can get around this problem if there is a subset of our observational data that is deeper than the bulk of our sample, and therefore of greater signal-to-noise, to a degree sufficient that it has negligible noise bias. This is often the case in real surveys that seek to detect high redshift supernovae, and it is approximated in the GREAT08 challenge with LowNoise_Blind data set (although the galaxy types in this set do not match those in the main sample exactly). In this case we do not calibrate each galaxy individually, but instead find a mean bias for the population and apply it en masse.

Biases $m(\theta)$ and $c(\theta)$ will afflict each of our galaxies, depending on the true galaxy properties θ . We can calculate mean values of this bias \hat{m} and \hat{c} , and apply these evenly to all the galaxies, such that the mean galaxy ellipticity (which is our goal) is correct.

The mean of the bias across the population is given by:

$$\hat{m} = \int m(\theta) p(\theta) d\theta. \quad (\text{E1})$$

We find the distribution $p(\theta)$ using fits to the deep data, and $m(\theta)$ using simulations - for each point θ_p in a grid in the parameter space we simulate many galaxies and determine the value $m(\theta_p)$. We can then interpolate between these values to do the integral, and finally apply the mean \hat{m} to all the galaxy estimates.

A similar process is used for the c bias, except that c varies directly with the PSF, so we use this information - we calculate the \hat{c} assuming a fiducial PSF with ellipticity e_0^{psf} aligned with the e_1 direction. The applied value to apply to each galaxy is then:

$$c_1 + ic_2 = \hat{c} \cdot \frac{e^{\text{psf}}}{e_0^{\text{psf}}} \cdot e^{i\theta_{\text{psf}}}. \quad (\text{E2})$$

¹³ Available at <http://www.ics.forth.gr/~lourakis/levmar/>

Table F1. Quality factor Q scores for GREAT08 from this work and the original challenge entrants.

Name	Branch								
	1	2	3	4	5	6	7	8	9
I3	817.2	1342.7	450.8	725.1	722.0	4702.5	84.6	382.7	99.8
I3-U	128.7	183.5	67.2	78.1	115.4	1417.1	25.3	2613.0	10.6
HB	256.4	654.4	241.9	184.2	158.9	225.3	176.0	272.1	131.7
AL	764.9	595.1	348.4	121.9	404.0	282.8	22.4	529.5	512.1
TK	224.7	569.6	737.4	161.9	238.5	63.6	54.3	142.9	59.9
CH	174.2	324.7	157.4	83.3	212.4	173.9	13.1	230.8	18.6
PG	21.7	26.5	26.3	37.8	44.2	28.5	50.6	23.1	82.4
MV	246.2	155.3	95.1	126.3	224.4	488.7	4.3	398.8	23.8
KK	222.2	554.8	448.6	89.8	230.4	100.0	3.1	163.4	33.7
HHS3	26.9	33.0	30.7	10.2	20.3	30.2	30.3	20.5	25.2
SB	161.4	93.4	226.3	114.5	155.2	19.8	4.9	26.0	8.6
HHS2	23.0	26.8	26.1	9.9	18.3	23.4	25.4	19.4	23.2
HHS1	67.2	109.0	22.3	43.0	60.9	232.7	8.8	1516.6	2.1
MJ	10.6	11.5	10.4	23.0	4.0	8.3	13.9	10.4	16.4
USQM	13.3	15.8	11.4	9.4	23.4	8.4	0.8	1.1	0.2

APPENDIX F: RESULTS TABLES

Tables F1-F4 give the GREAT08 Q , m , c_1 , and c_2 values for IM3SHAPE with (I3) and without (I3-U) noise bias calibration applied. Also shown are the scores for the other entrants at the time of the challenge, as described in Bridle et al. (2010).

Table F2. Multiplicative bias m scaled by 10^{-3} on GREAT08 from this work and the original challenge entrants.

Name	Branch								
	1	2	3	4	5	6	7	8	9
I3	15.3	3.3	-3.3	12.9	13.9	3.4	-35.8	-18.3	91.9
I3-U	37.6	25.4	49.8	35.1	36.2	8.9	49.0	3.3	116.0
HB	28.9	12.8	25.6	23.3	24.9	28.0	11.3	25.6	37.8
AL	9.9	-3.1	7.8	4.9	6.5	23.3	-87.5	-3.4	8.4
TK	26.6	14.3	-5.6	20.9	21.1	59.6	-57.0	37.0	-3.4
CH	17.8	7.9	11.0	11.7	16.7	9.5	-35.1	24.6	0.1
PG	70.3	58.1	58.9	70.5	67.0	38.3	-38.6	61.4	42.1
MV	-27.4	-35.8	-43.2	-36.6	-30.0	19.6	-188.4	-19.3	-84.4
KK	25.7	8.8	5.0	16.8	21.6	45.4	-254.4	27.6	-61.0
HHS3	74.1	64.1	67.0	70.2	69.8	66.2	63.8	87.6	74.5
SB	-36.9	-48.1	-22.6	-41.6	-35.6	-101.2	-167.4	-88.4	-7.3
HHS2	83.1	70.7	72.7	75.5	79.2	79.1	75.8	90.4	82.5
HHS1	54.5	39.1	89.6	43.7	46.6	-28.1	150.7	-8.9	306.0
MJ	61.1	47.1	58.0	54.9	53.3	-36.5	22.5	3.1	26.6
USQM	-52.5	-60.0	-62.7	-51.3	-49.2	-118.6	420.8	-393.5	882.2

Table F3. Additive bias c_1 scaled by 10^{-5} on GREAT08 from this work and the original challenge entrants.

Name	Branch								
	1	2	3	4	5	6	7	8	9
I3	-22.0	-31.8	63.8	-33.9	12.8	-17.8	-52.5	44.6	-254.3
I3-U	-67.0	-77.0	-94.6	-124.0	29.4	-29.0	-235.0	1.1	-304.0
HB	-4.4	-17.6	-16.8	-66.3	58.2	-45.2	59.4	-26.1	10.7
AL	-38.4	-51.5	-59.3	-118.6	34.8	-54.5	-79.9	-53.3	-51.3
TK	-26.3	-28.1	20.5	-63.5	9.5	-50.7	13.4	-44.4	154.6
CH	82.3	61.7	84.3	143.2	-22.0	93.1	-9.9	23.7	262.5
PG	188.4	177.9	182.1	-41.9	-87.5	220.0	-66.8	186.5	34.1
MV	8.5	-0.9	-37.5	-31.6	23.1	-6.7	-283.5	38.0	-79.3
KK	-49.2	-44.6	-47.1	-132.2	37.6	-57.8	-119.5	-62.8	-129.9
HHS3	-174.2	-165.6	-160.1	-387.2	75.8	-180.0	-168.8	-191.8	-196.5
SB	9.5	16.6	-54.5	18.7	16.0	135.4	-311.2	110.3	-459.1
HHS2	-177.6	-185.8	-176.0	-385.1	86.0	-196.5	-178.0	-195.3	-189.0
HHS1	-71.2	-74.8	-138.6	-169.4	56.4	48.4	-225.2	1.2	-431.7
MJ	-245.8	-249.7	-219.1	69.1	-696.9	-345.8	-92.1	-288.8	285.7
USQM	-138.0	-111.2	-114.0	-295.7	70.3	-118.8	-149.5	-161.5	-20.2

Table F4. Additive bias c_2 scaled by 10^{-5} on GREAT08 from this work and the original challenge entrants.

Name	Branch								
	1	2	3	4	5	6	7	8	9
I3	2.1	-10.6	-17.0	9.1	3.4	-0.6	9.4	-24.4	61.2
I3-U	18.5	5.6	41.7	42.1	47.9	3.5	75.8	-8.5	78.9
HB	15.3	11.6	31.2	43.2	28.1	19.0	31.6	15.3	-1.6
AL	19.6	11.0	39.7	49.6	49.3	15.8	-14.1	10.7	-1.0
TK	44.6	30.8	40.0	69.8	61.5	12.5	28.1	24.5	-92.7
CH	-16.2	-37.5	6.1	-39.4	-72.4	-33.0	53.4	-37.7	-24.0
PG	21.2	14.7	39.0	81.9	-9.0	11.0	79.9	5.0	40.3
MV	14.4	-14.0	26.9	27.7	-4.5	-11.1	82.0	-1.5	1.9
KK	18.7	1.7	41.8	60.1	39.4	9.6	10.2	19.6	-24.4
HHS3	54.9	63.1	83.7	148.3	193.6	57.3	70.9	69.6	56.5
SB	-15.1	-27.0	22.3	-32.8	-45.2	-53.7	87.5	-43.4	126.8
HHS2	69.4	65.4	84.5	150.6	185.9	67.3	58.1	69.3	48.3
HHS1	27.1	18.2	79.0	72.4	77.0	-20.1	49.2	-4.7	145.2
MJ	322.9	316.9	349.5	218.3	127.9	313.0	353.2	329.2	164.7
USQM	-114.4	-54.5	-182.7	-83.3	-33.3	22.0	-977.9	187.4	-1615.3